

# Bluespec-7: Scheduling & Rule Composition

Arvind  
Computer Science & Artificial Intelligence Lab  
Massachusetts Institute of Technology

February 28, 2007

<http://csg.csail.mit.edu/6.375/>

L10-1

## GAA Execution model

*Repeatedly:*

- ◆ Select a rule to execute
- ◆ Compute the state updates
- ◆ Make the state updates

Highly non-deterministic

User annotations can help in rule selection

Implementation concern: Schedule multiple rules concurrently without violating one-rule-at-a-time semantics

February 28, 2007

<http://csg.csail.mit.edu/6.375/>

L10-2

## Rule: As a State Transformer

A rule may be decomposed into two parts  $\pi(s)$  and  $\delta(s)$  such that

$$s_{next} = \text{if } \pi(s) \text{ then } \delta(s) \text{ else } s$$

$\pi(s)$  is the condition (predicate) of the rule, a.k.a. the "CAN\_FIRE" signal of the rule. (conjunction of explicit and implicit conditions)

$\delta(s)$  is the "state transformation" function, i.e., computes the next-state value in terms of the current state values.

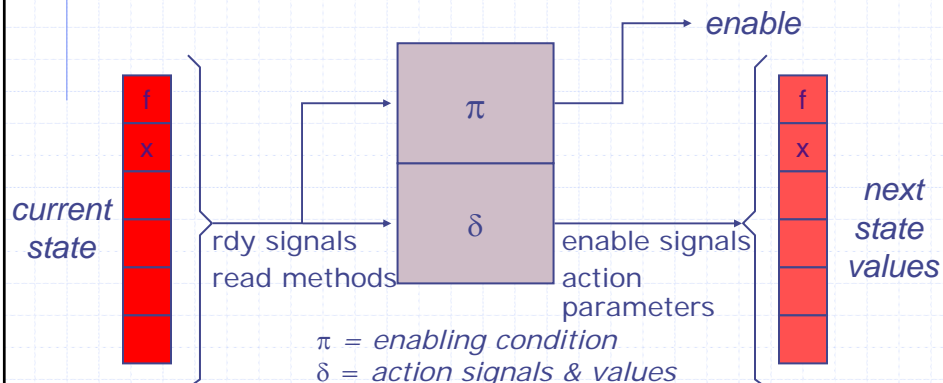
February 28, 2007

<http://csg.csail.mit.edu/6.375/>

L10-3

## Compiling a Rule

```
rule r (f.first() > 0) ;  
    x <= x + 1 ; f.deq () ;  
endrule
```



February 28, 2007


<http://csg.csail.mit.edu/6.375/>

L10-4

# Combining State Updates: *strawman*

$\pi$ 's from the rules  
that update  $R$

$\delta$ 's from the rules  
that update  $R$



February 28, 2007

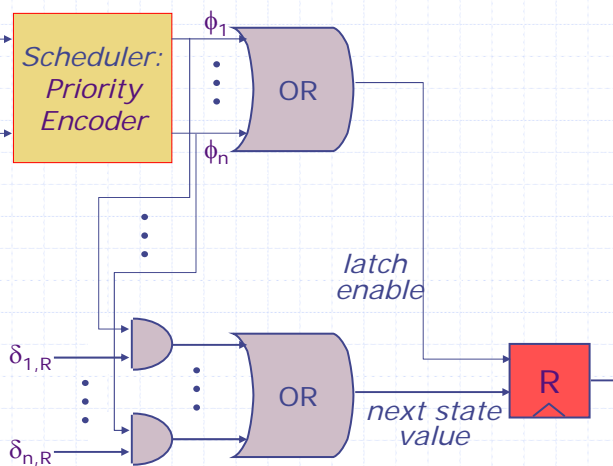
<http://csg.csail.mit.edu/6.375/>

L10-5

# Combining State Updates

$\pi$ 's from all  
the rules

$\delta$ 's from the rules  
that update  $R$




*Scheduler ensures that at most one  $\phi_i$  is true*

February 28, 2007

<http://csg.csail.mit.edu/6.375/>

L10-6

# One-rule-at-a-time Scheduler



1.  $\phi_i \Rightarrow \pi_i$
2.  $\pi_1 \vee \pi_2 \vee \dots \vee \pi_n \Rightarrow \phi_1 \vee \phi_2 \vee \dots \vee \phi_n$
3. One rewrite at a time  
i.e. at most one  $\phi_i$  is true

Very conservative way of guaranteeing correctness

February 28, 2007

<http://csg.csail.mit.edu/6.375/>

L10-7

## Executing Multiple Rules Per Cycle: *Conflict-free rules*

```
rule ra (z > 10);
  x <= x + 1;
endrule
```

```
rule rb (z > 20);
  y <= y + 2;
endrule
```

Parallel execution behaves like  $ra < rb = rb < ra$

Rule<sub>a</sub> and Rule<sub>b</sub> are **conflict-free** if

- $$\forall s. \pi_a(s) \wedge \pi_b(s) \Rightarrow \begin{array}{l} 1. \pi_a(\delta_b(s)) \wedge \pi_b(\delta_a(s)) \\ 2. \delta_a(\delta_b(s)) == \delta_b(\delta_a(s)) \end{array}$$

Parallel Execution can also be understood in terms of a composite rule

```
rule ra_rb((z>10)&&(z>20));
  x <= x+1; y <= y+2;
endrule
```

February 28, 2007

<http://csg.csail.mit.edu/6.375/>

L10-8

## Executing Multiple Rules Per Cycle: *Sequentially Composable rules*

```
rule ra (z > 10);
  x <= y + 1;
endrule
```

Parallel execution behaves like ra < rb

```
rule rb (z > 20);
  y <= y + 2;
endrule
```

Rule<sub>a</sub> and Rule<sub>b</sub> are **sequentially composable** if  
 $\forall s . \pi_a(s) \wedge \pi_b(s) \Rightarrow \pi_b(\delta_a(s))$

Parallel Execution can also be understood in terms of a composite rule

```
rule ra_rb((z>10)&&(z>20));
  x <= y+1; y <= y+2;
endrule
```

February 28, 2007

<http://csg.csail.mit.edu/6.375/>

L10-9

## Sequentially Composable rules ...

```
rule ra (z > 10);
  x <= 1;
endrule
```

Parallel execution can behave either like ra < rb or rb < ra but the two behaviors are not the same

```
rule rb (z > 20);
  x <= 2;
endrule
```

**Composite rules**

Behavior ra < rb

```
rule ra_rb(z>10 && z>20);
  x <= 2;
endrule
```

Behavior rb < ra

```
rule rb_ra(z>10 && z>20);
  x <= 1;
endrule
```

February 28, 2007

<http://csg.csail.mit.edu/6.375/>

L10-10

## Compiler determines if two rules can be executed in parallel

Rule<sub>a</sub> and Rule<sub>b</sub> are conflict-free if

$$\forall s . \pi_a(s) \wedge \pi_b(s) \Rightarrow$$

1.  $\pi_a(\delta_b(s)) \wedge \pi_b(\delta_a(s))$
2.  $\delta_a(\delta_b(s)) = \delta_b(\delta_a(s))$

$$\begin{aligned} D(Ra) \cap R(Rb) &= \phi \\ D(Rb) \cap R(Ra) &= \phi \\ R(Ra) \cap R(Rb) &= \phi \end{aligned}$$

Rule<sub>a</sub> and Rule<sub>b</sub> are sequentially composable if

$$\forall s . \pi_a(s) \wedge \pi_b(s) \Rightarrow \pi_b(\delta_a(s))$$

$$\begin{aligned} D(\pi_b) \cap R(Ra) \\ &= \phi \end{aligned}$$

These properties can be determined by examining the domains and ranges of the rules in a pairwise manner.

These conditions are sufficient but not necessary.  
Parallel execution of CF and SC rules does not increase the critical path delay

February 28, 2007

<http://csg.csail.mit.edu/6.375/>

L10-11

## Mutually Exclusive Rules

- ◆ Rule<sub>a</sub> and Rule<sub>b</sub> are mutually exclusive if they can never be enabled simultaneously

$$\forall s . \pi_a(s) \Rightarrow \sim \pi_b(s)$$

*Mutually-exclusive rules are Conflict-free even if they write the same state*

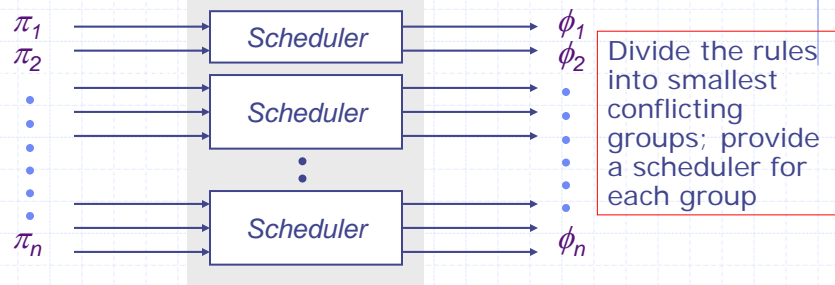
*Mutual-exclusive analysis brings down the cost of conflict-free analysis*

February 28, 2007

<http://csg.csail.mit.edu/6.375/>

L10-12

# Multiple-Rules-per-Cycle Scheduler



1.  $\phi_i \Rightarrow \pi_i$
2.  $\pi_1 \vee \pi_2 \vee \dots \vee \pi_n \Rightarrow \phi_1 \vee \phi_2 \vee \dots \vee \phi_n$
3. Multiple operations such that  $\phi_i \wedge \phi_j \Rightarrow R_i$  and  $R_j$  are conflict-free or sequentially composable

February 28, 2007


<http://csg.csail.mit.edu/6.375/>

L10-13

# Muxing structure

- ◆ Muxing logic requires determining for each register (action method) the rules that update it and under what conditions


Conflict Free (Mutually exclusive)



CF rules either do not update the same element or are ME

$$\pi_1 \rightarrow \sim\pi_2$$

Sequentially composable

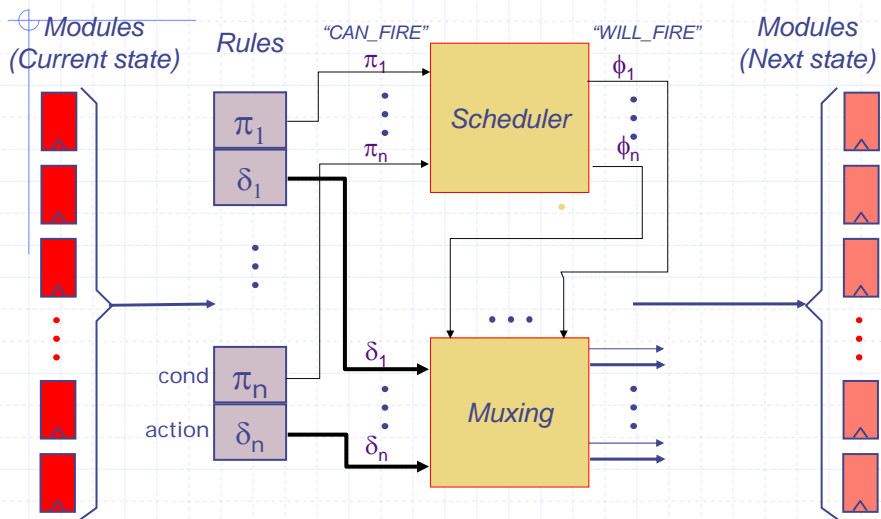


February 28, 2007

<http://csg.csail.mit.edu/6.375/>

L10-14

# Scheduling and control logic




February 28, 2007

<http://csg.csail.mit.edu/6.375/>

L10-15

## some insight Pictorially



- There are more intermediate states in the rule semantics (a state after each rule step)
- In the HW, states change only at clock edges


February 28, 2007

<http://csg.csail.mit.edu/6.375/>

L10-16



## Parallel execution reorders reads and writes




- In the rule semantics, each rule sees (reads) the effects (writes) of previous rules
- In the HW, rules only see the effects from previous clocks, and only affect subsequent clocks

February 28, 2007

<http://csg.csail.mit.edu/6.375/>

L10-17

## Correctness



- Rules are allowed to fire in parallel only if the net state change is equivalent to sequential rule execution (i.e., CF or SC)
- Consequence: the HW can never reach a state unexpected in the rule semantics

February 28, 2007

<http://csg.csail.mit.edu/6.375/>

L10-18

## Synthesis Summary


- ◆ Bluespec generates a *combinational hardware scheduler* allowing multiple enabled rules to execute in the same clock cycle
  - The hardware makes a rule-execution decision on every clock (i.e., it is not a static schedule)
  - Among those rules that CAN\_FIRE, only a subset WILL\_FIRE that is consistent with a Rule order
- ◆ Since multiple rules can write to a common piece of state, the compiler introduces appropriate muxing logic
- ◆ For proper pipelining, dead-cycle elimination and value forwarding, the user needs some understanding and control of scheduling

February 28, 2007

<http://csg.csail.mit.edu/6.375/>

L10-19

## Two-stage Pipeline



```
rule fetch_and_decode (!stallfunc(instr, bu));
    bu.enq(newIt(instr, rf));
    pc <= predIa;
endrule
```

Can these rules fire concurrently?


```
rule execute (True);
    case (it) matches
        tagged EAdd{dst:.rd,src1:.va,src2:.vb}: begin
            rf.upd(rd, va+vb); bu.deq(); end
        tagged EBz {cond:.cv,addr:.av}:
            if (cv == 0) then begin
                pc <= av; bu.clear(); end
            else bu.deq();
        tagged ELoad{dst:.rd,addr:.av}: begin
            rf.upd(rd, dMem.read(av)); bu.deq(); end
        tagged EStore{value:.vv,addr:.av}: begin
            dMem.write(av, vv); bu.deq(); end
    endcase endrule
```

February 28, 2007

<http://csg.csail.mit.edu/6.375/>

L10-20

## Two-stage Pipeline Analysis



1. fetch < execute
2. execute < fetch

February 28, 2007

<http://csg.csail.mit.edu/6.375/>

L10-21

## Scheduling expectations: execute < fetch schedule

```
rule fetch_and_decode (!stallfunc(instr, bu));
    bu.enq(newIt(instr, rf));
    pc <= predIa;
endrule
```

```
rule execute (True);
    case (it) matches
        tagged EAdd{dst:.rd,src1:.va,src2:.vb}: begin
            rf.upd(rd, va+vb); bu.deq(); end
        tagged EBz {cond:.cv,addr:.av}:
            if (cv == 0) then begin
                pc <= av; bu.clear(); end
            else bu.deq();
        tagged ELoad{dst:.rd,addr:.av}: begin
            rf.upd(rd, dMem.read(av)); bu.deq();
        tagged EStore{value:.vv,addr:.av}: begin
            dMem.write(av, vv); bu.deq(); end
    endcase endrule
```

February 28, 2007

<http://csg.csail.mit.edu/6.375/>

L10-22

## One Element FIFO Analysis

```
module mkFIFO1 (FIFO#(t));
  Reg#(t) data <- mkRegU();
  Reg#(Bool) full <- mkReg(False);      deq < enq ?
  method Action enq(t x) if (!full);
    full <= True;    data <= x;
  endmethod
  method Action deq() if (full);        first < deq ?
    full <= False;
  endmethod
  method t first() if (full);          first < enq ?
    return (data);
  endmethod
  method Action clear();
    full <= False;
  endmethod
endmodule
```

Expectation bu: (first<deq) < (find<enq)

February 28, 2007

<http://csg.csail.mit.edu/6.375/>

L10-23

## The good news ...

- ◆ It is always possible to transform your design to meet desired concurrency and functionality

February 28, 2007


<http://csg.csail.mit.edu/6.375/>

L10-24

# Register Interfaces

$read < write$

$write < read ?$



$read'$  – returns the current state when *write is not enabled*  
 $read'$  – returns the value being written if *write is enabled*

February 28, 2007


<http://csg.csail.mit.edu/6.375/>

L10-25

# Ephemeral History Register (EHR)

[MEMOCODE'04]

$read^0 < write^0 < read^1 < write^1 < \dots$



$write^{i+1}$  takes precedence over  $write^i$

February 28, 2007

<http://csg.csail.mit.edu/6.375/>

L10-26

## Transformation for Performance

```
rule fetch_and_decode (!stallfunc(instr, bu)1);
    bu.enq1(newIt(instr, rf));
    pc <= predIa;
endrule
```

```
rule execute (True);
    case (it) matches
        tagged EAdd{dst:.rd,src1:.va,src2:.vb}: begin
            rf.upd0(rd, va+vb); bu.deq0(); end
        tagged EBz {cond:.cv,addr:.av}:
            if (cv == 0) then begin
                pc <= av; bu.clear0(); end
            else bu.deq0();
        tagged ELoad{dst:.rd,addr:.av}: begin
            rf.upd0(rd, dMem.read(av)); bu.deq0(); end
        tagged EStore{value:.vv,addr:.av}: begin
            dMem.write(av, vv); bu.deq0(); end
    endcase endrule
```

February 28, 2007

<http://csg.csail.mit.edu/6.375/>

L10-27

## One Element FIFO *using EHRs*

```
module mkFIFO1 (FIFO#(t));
    EHReg2#(t) data <- mkEHReg2U();
    EHReg2#(Bool) full <- mkEHReg2(False);
    method Action enq0(t x) if (!full.read0);
        full.write0 <= True; data.write0 <= x;
    endmethod
    method Action deq0() if (full.read0);
        full.write0 <= False;
    endmethod
    method t first0() if (full.read0);
        return (data.read0);
    endmethod
    method Action clear0();
        full.write0 <= False;
    endmethod
endmodule
```

February 28, 2007

<http://csg.csail.mit.edu/6.375/>

L10-28

# After Renaming

- ◆ Things will work
  - both rules can fire concurrently

Programmer Specifies:

$$R_{\text{execute}} < R_{\text{fetch}}$$

Compiler Derives:

$$(\text{first}^0, \text{deq}^0) < (\text{find}^1, \text{deq}^1)$$

What if the programmer wrote this?

$$R_{\text{execute}} < R_{\text{execute}} < R_{\text{fetch}} < R_{\text{fetch}}$$

# Experiments in scheduling

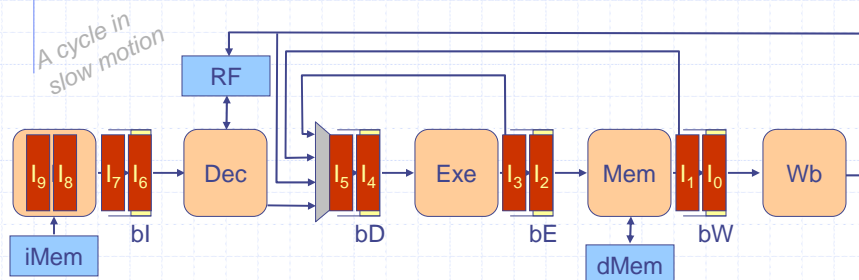
Dan Rosenband, ICCAD 2005

- ◆ What happens if the user specifies:

$$Wb < Wb < Mem < Mem < Exe < Exe < Dec < Dec < IF < IF$$

*No change in rules*

*a superscalar processor!*



*Executing 2 instructions per cycle requires more resources but is functionally equivalent to the original design*